RELEASE 1.0

8-94 *ESTHER DYSON'S MONTHLY REPORT* 3 AUGUST 1994

MULTI-USER VIRTUAL ENVIRONMENTS, PART II
by Jerry Michalski

Phil clicks on the doorway and passes into the garden that Zoe has designed. A few messages flash across the bottom of the screen as the software retrieves and assembles the different graphic elements of the scene. The garden's description is on Zoe's computer, but it points to components that exist elsewhere. The chairs and fence are generic. Phil's machine generates them locally and puts them in the right places. The gazebo is Zoe's special design, and transfers quickly. The nighttime sky looks lovely. It's an hour-old snapshot straight from an observatory. Next year she'll use a video backdrop of a daytime autumn scene. Zoe's friend Juan designed the pond, with fish bred by El-Fish; it's on its way from his machine.

Zoe has licensed special texture maps from a third party. Those arrive, and the scene is rendered within a couple of seconds. But wait. The birds she selected appear to be very popular: Their server is bogged down...now it's crashed. Looks like there will be generic birds in this scene.

In large volumes, activities like this one generate transactions that strain every part of a networked computing and communication system. If you look at these activities as a curious experiment in fictional 3D spaces, they may not seem important.

But it's possible that the flat, document/desktop-metaphor world we are familiar with so far is a warmup exercise for more complex and useful virtual worlds to come that will find common use in business. Perhaps our future computer file systems will be abstract, semantically arranged dataspaces or resemble our real-world offices; our business meetings will be held in virtual meeting rooms with avatars sitting in software armchairs; or our interactive tvs will have virtual-world interfaces.

Whether or not any of these interfaces becomes commonplace, their underlying infrastructures are precursors of architectures to come. To support these highly distributed and dynamic 3D environments, developers have to rethink client/server architectures and devise ways to create and share objects that exist in many places and will be viewed from many perspectives. To make these spaces friendly and useful enough that people are drawn to them out of more than mere curiosity will require advances in user-interface design and navigational aids. They may develop visual ways to communicate context ⟹

WELCOME CHRISTINA AND EDVENTURE.COM!

INSIDE

VIRTUAL ISSUES 1
Of ring worlds and 3D.
Wanted: virtual semiotics.
VISITING THE ARCHITECTS 6
daVinci Time & Space.
Xerox PARC's Jupiter and Agorics.
VRML and Labyrinth.
RESOURCES & CALENDAR 18
Names, phones, dates.

and meaning. To make the whole thing a viable business proposition will require solid transaction underpinnings and global scalability.

Last month we described some virtual spaces, most of which run on centralized servers. We began with Apple's QuickTime VR, which allows one person to move through space defined on a single host computer, and progressed to systems that support multiple people on the same server or several linked servers, such as Virtual Universe and Kesmai's Aries.

This issue of **Release 1.0** takes another step in the same direction: Most of the companies we cover are building virtual places that can hold large numbers of people and are hosted on many distributed servers. In some cases, the participants can create their own objects or places to share; in others, the environment is a context for shared activities.

Thorny (but fun) issues remain

Exploring virtuality leads developers to address a series of interesting questions, many of which are simultaneously practical and philosophical. For example, what kind of space should they build? Euclidean? Cartesian?[1] Hyperbolic (negatively curved)? Toroidal or ring-shaped? A Dyson Sphere? Conceptual or semantic? Some combination?

Individuals could have their own views on the space. For the venturesome or adventurous, a 3D world could have wormholes that teleport occupants to other spots, or doorways to worlds with other spatial models, where different rules of physics apply.

A compelling reason for agreeing on a Cartesian 3D space is that it gives people a way to wander about and meet other people. If there's no agreement on where the village square is, people can't stroll there and run into each other. Of course, a 3D virtual world will likely have real-estate battles, just as the real world does, despite the fact that virtual spaces can be collapsed inside each other. If people's sense of orientation depends on there being a point or axis of origin, then space nearest those foci will be dear. (Who gets to own 0,0,0?) There are only so many pixels to work with on a given screen.

Non-3D space is not immune from land-rush phenomena: Witness the current race to register domain names on the Internet, a race that is reminiscent of but less restrictive than those for vanity license plates or 800 numbers. (Who will register AAA Internet Services?)

Developers of virtual environments also have more tactical considerations, such as whether the space should be 2D, 3D or somewhere in between. Two-and-a-half dimensional space could have flat characters in a 3D space, or vice-versa. Avoiding full 3D will probably reduce the processing load on client computers.

---

1 *Euclidean space follows Euclid's axioms and definitions about elements such as straight and parallel lines, perpendicularity and the angles of triangles (ca. 300 BC). In the 17th century, French philosopher Rene Descartes added an origin point (among other things) and created the familiar Cartesian coordinate system.*

Spatial integrity: A place to hang your hat

Choosing a model for virtual space involves the familiar tradeoff between usefulness and power. On the one hand, the space should be intuitively navigable and offer a useful frame of reference for occupants. Real space is rich with artifacts and features that people can use to know where they are and what they can do next, such as store-front windows on a city's main street or mailboxes on a country lane.

On the other hand, applying the physical world as a model may obscure or preclude many powerful features of virtual space that we would otherwise happily adopt. We may imagine things as opaque when in fact they are transparent; we may think they are distant when they are really close at hand. We may limit our thinking by imposing the wrong mental constructs.

For example, one of the fascinating properties of virtual space is that distance collapses. When you click on a tagged item in Mosaic (the World Wide Web[2] browser), you are transported to your destination as quickly as network traffic permits. There is no sense of travel. You don't climb into a vehicle and wait to get there, while reading billboards. Of course, there are some people who would like you to sit still for a few ads and may help you pay for your time in cyberspace for that privilege. Other kinds of spaces probably have similar benefits and drawbacks.

Wanted: semiotics for virtual environments

If the designers choose a semantic space, what is descriptive of a data structure beyond its filename? Different people might need very different frames of reference for their activities. The sales manager of a transportation company may see the world in terms of maps and schedules; the human resources/facilities manager may think in terms of office layouts, organization charts and personnel files; a chemical plant supervisor may visualize process diagrams and flow charts.

In any kind of space, people need cues about what's there and what they can do with it, cues that human-factors expert and Apple fellow Donald Norman calls "affordances." People need new road signs for the new environment and objects that broadcast their uses. This goes beyond icon and toolbar design to how shape, color, location and relationship give meaning. There's plenty of room for fruitful exploration of these topics now that there's enough processing power and network connectivity to make the problem more than a thought experiment.

You say "to-may-toe," I say "to-mah-toe"...

How can many people share highly distributed space and objects in a way that is agreeable to all of them? In centralized systems, such as Kesmai's Air Warrior, all participants get scene and object descriptions from the same place. Distributed architectures will likely lead to the creation of consensual spaces: Participants may not have the same representations for objects and may not even see (or hear) the same objects.

---

2 *For more information on the World Wide Web, see Release 1.0, 1-94.*

Say Zoe has chosen to see all chairs as colorful, plush, overstuffed, animated seats -- along the lines of "Chairy" in Pee-Wee's Playhouse -- and Phil has chosen a sober, distinguished Frank Lloyd Wright style. When Phil's avatar sits in Zoe's chair, she may laugh, but Phil may find her reaction inappropriate because his context is different. Phil may appear to walk through Zoe's broader chair. Items of different sizes will attract attention differently. Perhaps if Zoe chose a conceptual view and Phil a physical one, they would have an easier time communicating, since they would be less confused about the fine points.

Of course, there would be no challenge if there weren't some evil involved. Some of the people moving about in these spaces may be up to no good, so the protocols and services have to assume the worst of all participants. That's the idea behind the script interpreters in Safe-Tcl and General Magic's Telescript (see **Release 1.0**, 2-94). Robust multi-user virtual environments are likely to need similar protection.

Sci-fi roots

When people start to design virtual worlds, all the issues of real worlds emerge, multiplied by the alternate worlds that the new medium allows them to construct. Science-fiction writers have been building alternate worlds for quite some time, and the architects of virtual space draw inspiration from them.

The world is holographic and recursive. Cyberspace will contain not only reflections of our world in all its complexity, but also manifestations of all the potential and imagined ones that people care to create. Reality and virtuality will point to and wind around each other.

Science fiction and science-related non-fiction are much closer to the action than ever before. There is a productive interplay between technological developments and description of those developments, both real and fantasized. Increased bandwidth and processor power offer designers more opportunity to create imagined realms.

The current synthesis among once-separate disciplines is particularly fascinating. Science fiction, design, philosophy, physics and computer science are already intertwined: now in the fusion furnace of cyberstuff they are melding with social sciences, community building and architecture, to name a few disciplines. One recurring influencer is architect Christopher Alexander, whose work blends many of these perspectives (see box).

The ubiquitous Christopher Alexander

Although cleverly disguised as architectural theory, the work of UC Berkeley professor Christopher Alexander applies to a far broader realm of activity. Essentially, his work is about the collaborative design of spaces and places by their eventual occupants, aided by gentle guidelines -- think of them as nuggets of abstracted wisdom. Alexander calls the collected guidelines a pattern language, since they communicate essential patterns of living that translate into good design principles.

A Pattern Language is also the name of one of his books (Resources, page 18). It describes 253 interrelated and interacting patterns that relate to the design of place, from activity nodes and connected play to shielded parking, entrance transitions and the flow from room to room. (Its rich web of relationships makes it an especially good candidate document for the World Wide Web.)

Equipped with this pattern language, ordinary folks can create better spaces for themselves than if left to their own devices. For example, a guideline for a room to feel warm and inviting is that it have natural light from at least two walls. While this is not a feature an amateur architect might glean, it is memorable and compelling, and the amateur is able to use it. That's the sign of a good guideline.

Organic, decentralized design

In Alexander's conception of architecture, the design emerges from the collective instincts, experience and desires of those it most affects. The underlying principle is to loosen control. It has an organic growth model, rather like the Internet's or like work on artificial life and related subjects at the Santa Fe Institute (see **Release 1.0**, 4-92 and 5-92). In 1974, Alexander started the Center for Environmental Structure, where he and five colleagues help create places for people to live and work.

Write to us at EDventure.com!

We now have our own domain name on the Internet (but not our own host machine...not yet). See the masthead, page 18, for our new e-mail addresses. In addition, with the WELL's help, we've put the June 1994 issue on the World Wide Web. To visit it, point your Web client to <http://www.well.com/Community/Release1/>. We will place issues of **Release 1.0** in the public space from time to time, as a way to contribute to the community discussion when the topic and the setting are particularly appropriate.

VISITING THE ARCHITECTS

The companies covered in this issue clearly have not resolved all of the issues, but they are considering most of them in their designs. Design intent is crucially important in setting architectural direction. It's the difference between a system that needs to be patched and rewired and one that is gracefully extensible, even though it may not incorporate all of the futuristic aspects early on.

We cover three efforts to design virtual environments: daVinci Time & Space, a startup designing an environment for kids; Xerox PARC's Jupiter project, a virtual space for professionals; and the Virtual Reality Markup Language mailing list, an open, collaborative effort to create a public specification as a starting point for distributed VR on the Internet.

DAVINCI TIME & SPACE: NEW DIMENSIONS OF PLAY FOR KIDS

New-media companies are sprouting everywhere; the market for interactive, multimedia software for kids is particularly crowded. But few of the new entrants have the vision or the nerve of daVinci Time & Space (dVT&S), a San Mateo, CA, startup that intends to create safe and compelling virtual places for kids ages 3 to 12 where they can not only go on adventures and interact with software objects but also meet and interact with each other. Depending how the system is designed and what kind of infrastructure is available, the kids might be able to see and hear each other, too. Although you could call this interactive tv, it offers much more.

Imagine a virtual playground where kids from different cities can meet and play some software-enhanced variant of tag or use their imaginations to turn the monkey-bar set into a zoo or submarine. A third-party developer might create a great submarine simulation, available as a "ride" at extra cost. Another could offer a room where kids could play in a weightless environment; a third could offer an Amazon adventure (above water or below).

Co-founders Carol Peters and Jeff Apple think of dVT&S as being built atop four industries. It draws on the plot, characters and high production values of movies and television; the reach, infrastructure and bandwidth of cable tv; the software innovations, interactivity and aesthetics of multimedia CD-ROMs; and the social nature and many-to-many connectivity of online services.

Put more concretely (and hypothetically, since dVT&S isn't talking much publicly now), the environment they seek to build might blend the social online games of the ImagiNation Network with the persistent characters, storytelling and educational content of Sesame Street. It would hopefully be as palpable and engrossing (and multi-user) as your neighborhood Discovery Zone or McDonald's playground, but you can reconfigure the toys, and there are optional plots and adventures, as well as those persistent characters. It also has advertising -- but hopefully a more subtle kind than the way Prodigy started. More on the advertising in a moment.

All of this will require some as-yet-unbuilt high-speed infrastructure. Today's telephone, cable tv and wireless systems aren't equipped to host such services (see **Release 1.0**, 11-93). In fact, most of the 500-channel-

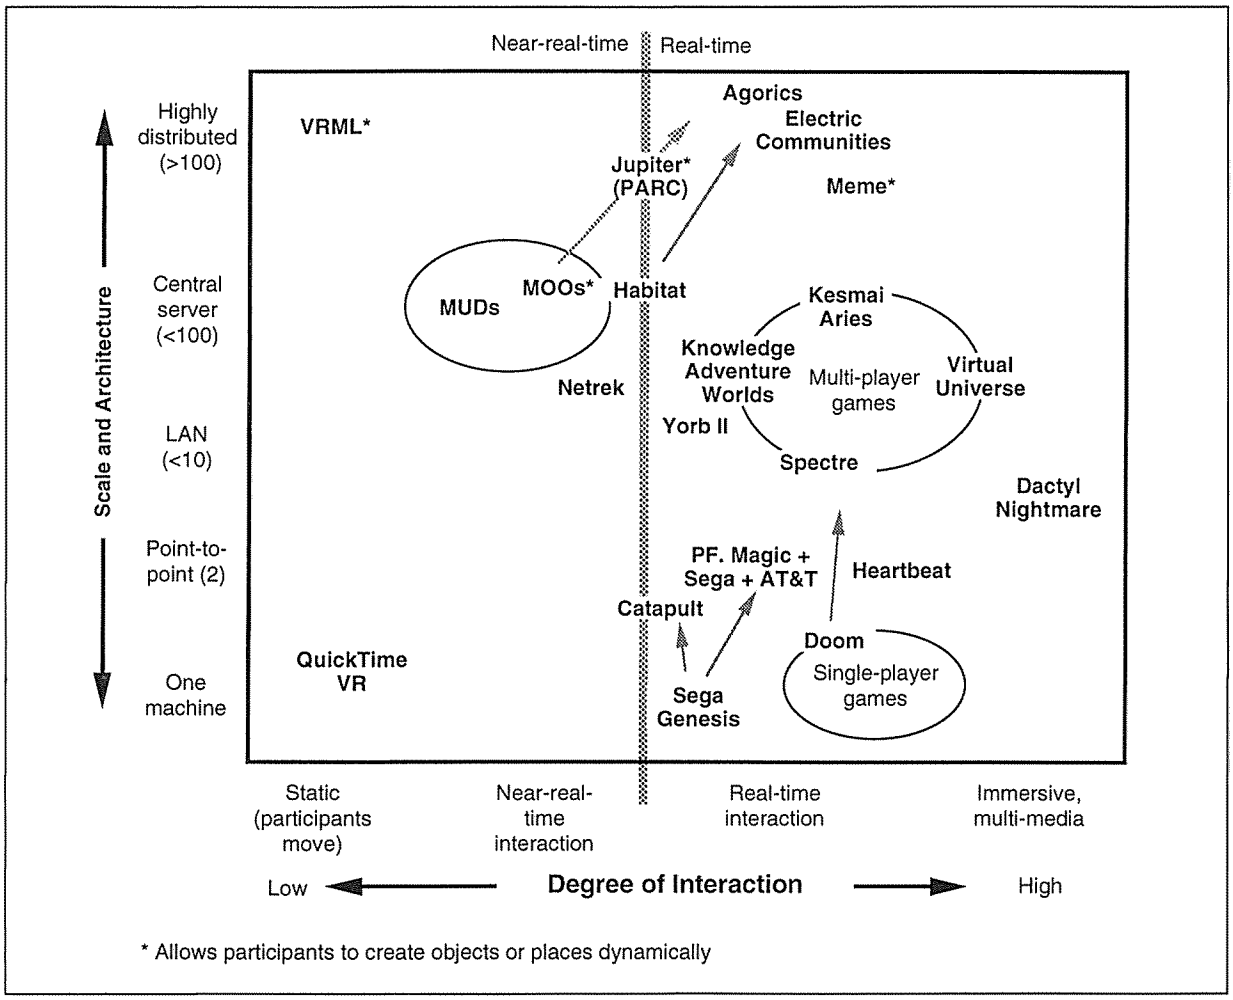

This diagram maps the companies covered in this and previous issues, using some of the dimensions described last month, such as how the virtual environment engages the senses (text, rendered graphics or immersive VR); how many people can participate at one time; how distributed the system is; and how dynamic it is (pan, zoom and move vs. create, share and sell novel objects).

There's considerable movement over time. For example, many single-player games go multi-player; some get live audio. MUDs and MOOs lie to the left of the real-time boundary because they process events sequentially. The Jupiter project adds real-time audio and slow-frame video, which brings them across the threshhold. As Jupiter moves to a more distributed engine, its interactivity will increase, as will the number of participants it can support.

Habitat originally ran on Commodore 64s with 300-baud modems (see **Release 1.0**, 7-93). Later this year it will be reinstantiated in a greatly improved form under a different name, with client software for Macs and Windows machines.

future trials in progress or in plans today aren't adequate: They focus narrowly on digital video-on-demand. To mitigate risk, dVT&S's engineers are designing the system to support a variety of infrastructures.

The more capable systems may be the platforms for early, limited-function pilots of dVT&S's wares in late 1994, with larger system trials in 1995. Luckily, the company's blue-chip investors -- Oak Investment Partners, Greylock Management and Venrock Associates -- consider dVT&S a long-term investment. So far they have put up $5 million; dVT&S estimates it will take between $25 and $30 million to launch the service.

Combined talents

Carol Peters and Jeff Apple founded dVT&S in August 1993. Peters, who studied architecture at Harvard, spent 16 years at DEC. There she rose from an entry-level job to running hardware development teams, culminating with the design of DEC's first RISC entry, the DECstation 3100. Then she went to Silicon Graphics, where she ran the Iris Indigo workstation project. The powerful, flexible and affordable Indigo changed the dynamics of the multimedia industry. Having designed the computer she really wanted to build, Peters left SGI late in 1992 to pursue her own interests, which drew her back to her roots in design -- only this time for kids.

Peters met co-ceo Jeff Apple, who has produced over 1,500 commercials and a few movies (including *In the Line of Fire*), through a mutual friend. Each wanted to move toward the others' industry, they had similar ideas about what was needed and they hit it off, so they decided to get funding and start daVinci Time & Space. Apple and a few other staff members work in Los Angeles. The 15 employees bring experience from jobs at Disney Software, Kaleida, Ketchum Advertising, McKinsey, the MIT Media Lab, the tv show "Pee-Wee's Playhouse," Q2 (QVC's second channel), The Shopping Channel and Xerox PARC.

Stay tuned, folks!

As dVT&S's staff creates its service platform, it will have to resolve two sensitive issues: How to ensure that third parties create compelling content with high production values, while preserving the overall tone and quality of the experience, and how to create advertising that is appropriate to the medium and its audience.

Other industries manage the issue of consistency reasonably well, though an interactive software service is quite a different animal. Stylebooks help magazines and serialized novels maintain consistency through many issues and with multiple writers; Disney maintains consistent style and behavior for the same characters, even though they are drawn by many people. DaVinci needs to create a software authoring environment that enforces style consistency and quality, but doesn't stifle creativity.

This game of tag was brought to you by...

Although daVinci Time & Space has made clear that its virtual environments will contain advertising, the fact that children will be the primary occupants makes the way the advertising is presented a particularly sensitive issue. Early interactive-ad concepts were of billboards that do something

when kids pass or point to them; dVT&S claims to have moved beyond that, but is keeping quiet about it (see **Release 1.0**, 12-93).

In a social medium, companies can gain beneficial exposure by engaging in appropriate social activities. Companies may sponsor activities or offer services, the way McDonald's puts small playgrounds in some of its stores and sponsors Ronald McDonald Houses near hospitals. Or they may help subsidize specific audiences whose access to the medium is limited.

Martina's not the only superstar who can wear product endorsements or do testimonials: Imagine a computer-generated Snow White with "Tide" embossed subtly on her blouse, or Mickey replacing Elsie the Borden cow with a first-hand consumer perspective on cheese -- rather than the producer's silent stare. Wait; we're kidding!

JUPITER: BEYOND LAMBDAMOO

In 1990, Pavel Curtis started the LambdaMOO at Xerox PARC.[3] LambdaMOO has since grown to around 8,000 occupants -- roughly the size of the WELL, a mid-sized regional online conferencing system -- and is considered to be the axis MOOndi. If you want to start a MOO of your own, you copy the files from PARC's server.

Now Curtis and his colleagues David Nichols and Mike Dixon are taking LambdaMOO beyond text. Rather than evolve it toward a highly visual and animated 3D world, the Jupiter team has drawn on what works well in LambdaMOO (the sense of place and presence, the persistent objects) and added features that seem useful, so they can experiment with them. In fact, the Jupiter team consciously stays close to the conventional GUI interface so that there isn't much of a leap to make from ordinary work with documents and files to the social environment of Jupiter. The idea is that people will keep Jupiter open on their screens all day.

Although that means Jupiter is not as visually striking as the systems described elsewhere in this and the June issues of **Release 1.0**, Jupiter may be more relevant to business users, which is Xerox's goal. Jupiter isn't a classical, short-term, product-development effort: It's an internal initiative to explore new work environments that support communities of practice. It may lead only to streamlined operations within Xerox, or possibly to partnerships to develop new businesses.

---

3 *We introduced LambdaMOO briefly in Release 1.0, 6-93 and 6-94. MOOs are textual, multi-player environments that allow participants to modify or create objects, characters and places. For a further description of MOOs and MUDs, see those issues, as well as 7-93. MUDs are multi-user dungeons; MOOs are MUDs that are object-oriented.*

What's new

Jupiter's developers have left the underlying LambdaMOO engine intact. They will leave it behind later when they move to a system that has global scale and robust security.

A peek at Jupiter

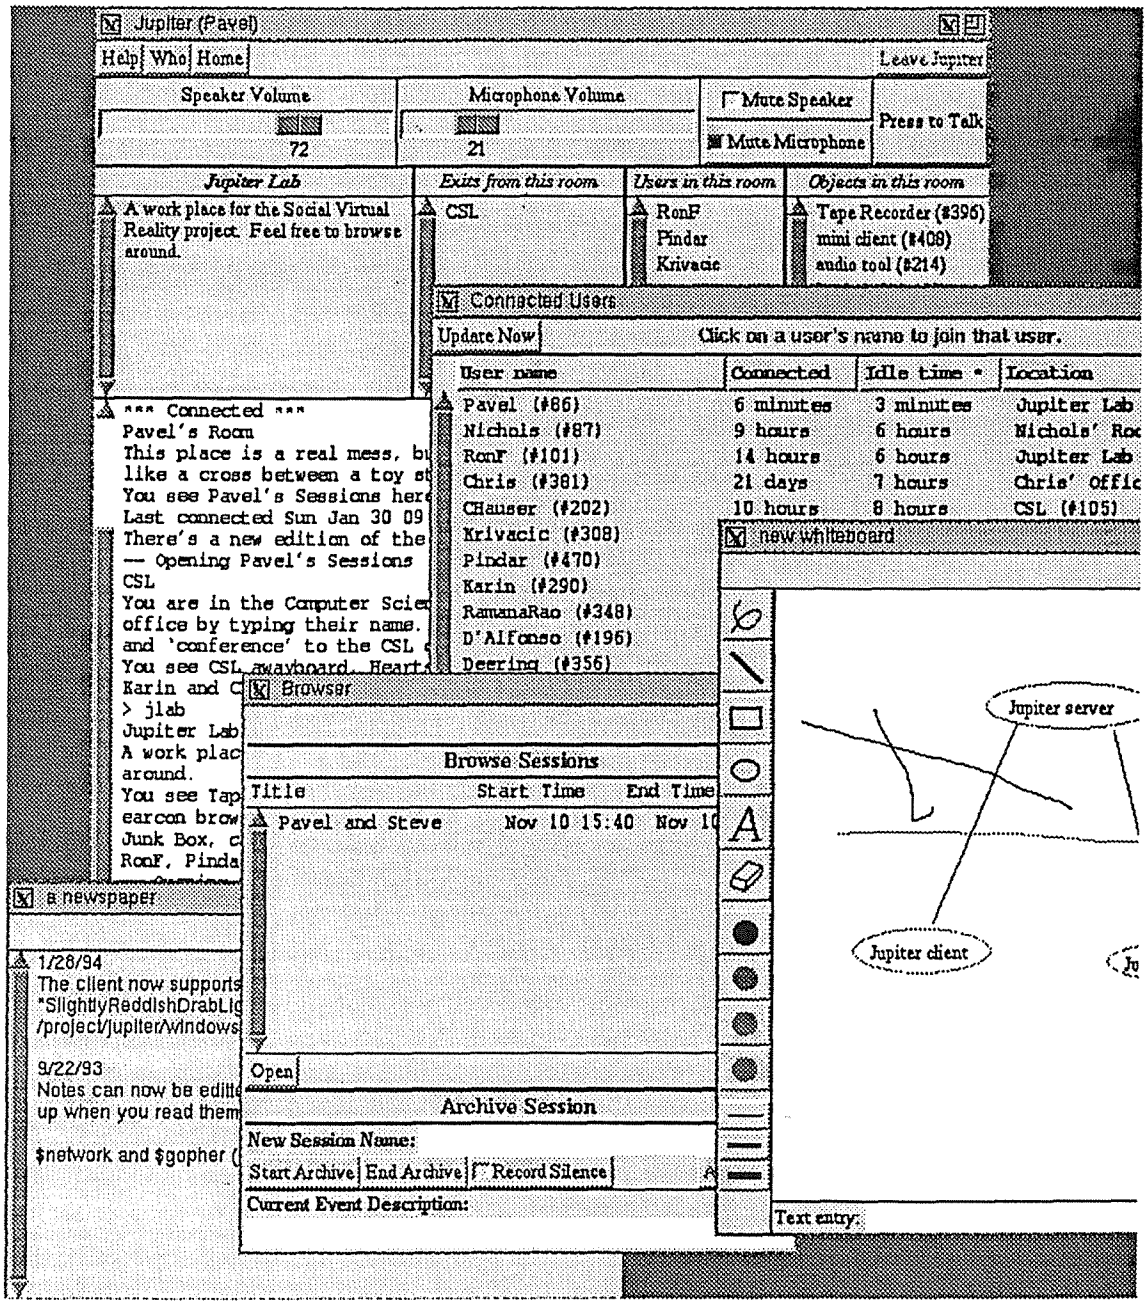

The developers have focused on improving the user interface, adding real-time communications capabilities and improving participants' ability to create collaborative-work tools. Jupiter's user interface (screen shot, opposite page) is much simpler to use than a basic MOO's. Instead of forcing participants to remember, then type, "open door", "e" (shorthand for "go East"), or "@examine book" (to see what a nearby book is or contains), Jupiter's interface presents that information in different panes of a window (the topmost window in the screen shot).

Users can easily tell who is around, what the obvious exits are and what objects are around for inspection and use. They can point and click to inspect and navigate the MOO, which is quick and requires less thinking about the medium -- a great improvement over the plain-text interface. This interface is specific to Jupiter.

Communications features

The graphical user interface makes existing LambdaMOO features easier to use, but the Jupiter team has also added new features, most notably live audio and video, which appear in a separate window.[4] You can press a button to speak to everyone in the virtual room (press-to-talk mode), leave the microphone on all the time or "whisper" to one occupant. Whispers are preceded by a sound like the chirp of a Star Trek communicator.

Jupiter can show several video feeds simultaneously. The communication window has a button labeled for every occupant of the "room" you're in, and shows live video images of the people you have chosen to see. To change whose video you see, drag one of the other buttons over an existing video window, and the two trade places.

Because the LambdaMOO software is not a real-time communications engine or a switch, the audio and video run parallel to the LambdaMOO but are invoked from its interface. When a user navigates the text window to enter a room, Jupiter can set up (and tear down) the video and audio sessions automatically. Perhaps we will one day have several Jupiter "rooms" open on our desk, one for each community or project team we belong to.

About 30 PARC people are now testing Jupiter, which is still under construction; some telecommute from home with Jupiter over ISDN. Astrophysicists at NASA/JPL and other sites are installing Astro-VR, a version of Jupiter without video. They plan to use it for distance conferencing, sharing astronomical images and private e-mail and bulletin-board features.

---

4 *Jupiter users have microphones and (optional) cameras mounted near their computers. The video refreshes at only five frames per second, has medium resolution and doesn't synchronize with the audio feed, but it is good enough to show whether others are present, have visitors or are absorbed in activity. Besides, several of these videos can be running in one palette at the same time (using the Internet multicast protocols) -- all on standard Ethernet.*

Hanging out in the MOO

Hanging out informally in virtual spaces is central to Jupiter's purpose. Phone calls and e-mail messages are useful when people know they want to get something done. Jupiter is designed to explore and enable subtler forms of interaction. It seeks to support communities of practice[5] with a fluid medium within which participants can be aware of each other's presence without feeling threatened or intruded upon, and in which they can shift easily between states of fruitful collaboration, observation and relaxation, of focus and play. Designing such a system requires a delicate balance between explicitness and implicitness, as well as careful attention to the border area between a group's center and its periphery (see box).

Audio delivers much of the value. Curtis says, "The use of audio is clearly important, because casual interactions must be essentially effortless, and speaking is easier than typing." Jupiter users typically leave the speaker on and use the microphone when they want to speak. The team would like to implement a feature that filters some audio from "adjacent" rooms to users, to give them some peripheral awareness of nearby activities.

Border issues and legitimate peripheral participation

Jupiter is part of current research at Xerox PARC into how technology can help people build and support communities of practice, which John Seely Brown, PARC's director, described at this year's PC Forum (see **Release 1.0**, 3-94). Current work focuses on the value of participants who are not in the center of a group's activities, and pays special attention to the place where the center and the periphery meet, which Brown calls the border.

Call screening is a great example of the border area: The called party can choose whether or not to engage the caller (stay in the periphery or step into the center) without being explicitly rude. The example Brown used in his speech was of Xerox copier field-service reps who were armed with two-way radios that they could keep on much of the time, if they so desired.

By listening to each other as they puzzled through a nasty diagnostic session, reps with experience could jump in to help build a plausible narrative and possible solutions. Others could listen and learn. This is an example of legitimate peripheral participation. People at the border are not responsible for events in the center, but they can participate if they want to.

Jupiter offers participants ways to communicate the desire to interact as well as the wish for privacy. Some rooms in the MOO are casual. Gardens, lounges or cafes may have games lying about, or newspapers (assembled auto-

---

5 *The usage of the terms "communities of practice," "periphery" and "legitimate peripheral participation" at PARC derives from studies of learning by Jean Lave and Etienne Wenger in 1991 (see Resources, page 18).*

matically from newswires and community contributions, of course). Being in them is an implicit invitation to interaction. On the other hand, a closed door is often a request for privacy.

Degrees of privacy could be manifest in various ways, ranging from a physical sensor on your office door that is reflected somehow on-screen to a privacy "knob" on the interface. You could mix the virtual rooms with do-not-disturb settings, effectively creating content filters for inbound messages and for people who want to communicate. Phil could be available in the room related to his deadline, but not available elsewhere. Jupiter has more ways to indicate availability than a desire for privacy.

Who's there?

The Jupiter team's exploration of ways for people to indicate their presence, trying to avoid sounds and events that confuse or overwhelm users. There's a fine line between useful information and intrusion. The physical and virtual worlds can work together: Many PARC employees wear infrared tags that allow the building to track them. With proper authority, people can get notified when others move about.

Participants using the video element in Jupiter are more visible than they would be even in close quarters in a small office. In Jupiter, you can see what each participant is doing, unless he or she decides to draw virtual blinds. In real life, people would probably face away from each other or use partitions.

Persistent objects and tools

Finally, Jupiter users can leave files, objects or programs in the environment, using a scripting language called MOO. All objects they create are inherently sharable without special coding. Dixon created a shared whiteboard in 28 lines of MOO code.

With this capability, scientists can leave specific programs in the MOO that colleagues can send their data files to for processing. They can "hang" images or charts on the walls of their virtual lab. The tools can access other computers, such as an information-retrieval interface that drives a Dialog database search now in use in the MOO. Better still, research librarians, network troubleshooters and other skilled people can help participants learn to use the tools appropriately by using Jupiter's communication features.

MOOving forward

There is a good reason that the modifications that the Jupiter team has made to LambdaMOO are mostly in the client software or in systems that run parallel to the MOO server: The LambdaMOO server code doesn't scale well and will be left behind eventually. In their quest for a secure, global infrastructure for Jupiter, the team has been heavily influenced by Joule, a language for creating robust, secure and scalable infrastructures for multi-user systems, from startup Agorics (see box, page 14).

There are many objects that the Jupiter team might build later on. Imagine some telephony integration: When your real phone rings, you could answer

Agorics' distributed-system platform

The agora is the Greek open square or marketplace. It symbolizes the way Agorics addresses the problems that choke most distributed systems: concurrency control, deadlocks, security and resource management. Agorics' system manages resources with economic markets, including auctions, price feedback, supply and demand. Participants can own, buy and sell resources. There's a mechanism everyone can use to allocate resources that allows individual entities to set up their own policies.

In a joint project with Sun Microsystems, Agorics developed a computational market for allocating the bandwidth of a fiber-optic ATM network. Agorics demonstrated a scalable digital video application built using this market in the Sun booth at Interop last May.

Agorics' Joule language addresses the problems of transparent distribution, concurrency, deadlock, security and resource management. Agorics is developing Joule in cooperation with Xerox PARC and Electric Communities (see **Release 1.0**, 7-93). The prototype implementation is being built in Smalltalk. As we just mentioned, Xerox PARC is considering Joule as a platform for a highly distributed MOO. Electric Communities will use Joule as the foundation for a global cyberspace project.

Agorics' president is Ann Hardy, formerly president of Key Logic (see **Release 1.0**, 5-88) and a vice president at Tymshare for over ten years. Agorics' architects are Dean Tribble, formerly with Xerox PARC, Xanadu (the hypertext project based on Ted Nelson's ideas, **Release 1.0**, 7-89) and Apple; Mark Miller (Xerox PARC, Xanadu); and Norman Hardy, a computer and communications security expert also formerly with Key Logic and Tymshare.

it with your Jupiter interface and conference in any other participants. You could then move within Jupiter, and share the call with others.

Lead, follow or get out of the way

Jupiter redefines communications within traditional markets and across markets. It offers groupware/collaborative-work features not available in e-mail and shared-document database systems such as Lotus Notes. Jupiter's user interface may prove to be a more intuitive way than dialog boxes to express important things such as availability. The interface also offers insights into computer-telephone integration and call control. Finally, as the MOO engine is improved, it may help define important distributed-system infrastructure concepts.

The Jupiter project shows a more measured and foresighted attitude toward communications than even most communications companies have. Although the project's objectives are loose and inward-focused for the moment, Jupiter offers Xerox plenty of opportunities to create new businesses or partnerships. Jupiter might be just the elixir that Xerox needs to rejuvenate its operations. Others may find it stimulating as well.

VRML, THE VIRTUAL REALITY MARKUP LANGUAGE

One of the popular topics at the recent World Wide Web conference in Geneva was how to implement distributed virtual reality on the Web. Mark Pesce and Brian Behlendorf, who didn't know each other before the conference, want to harness that interest to accelerate the development and acceptance of an open, platform-independent standard for Web-based VR. To that end, they have started a mailing list[6] for the Virtual Reality Markup Language (VRML)[7] and put some information about VRML and the list participants on the World Wide Web.

Early discussion on the VRML list ranged over topics as broad as the questions in the introduction to this issue of **Release 1.0**. Behlendorf and Pesce have kept the list focused so it can have concrete results, which others can then experiment with more broadly. In that spirit, the VRML list participants are now reviewing 3D scene- and object-description languages. Candidates include:

- **CDF**, the Cyberspace Description Format, developed by Autodesk.
- **FFIVW**, the File Format for the Interchange of Virtual Worlds, by Bernie Roehl and Kerry Bonin, is from the computer-games industry. The name speaks for itself.
- **Meme**, Immersive System's Multitasking Extensible Messaging Environment, proposed by Marc de Groot, president of Immersive, and described in **Release 1.0**, 6-94.
- **MSDL**, the Manchester (UK) Scene Description Language, derived from the next candidate, OOGL.
- **OOGL**, the Object-Oriented Graphics Language, developed at the Geometry Center in Minnesota and popular as a geometric visualization format in mathematics, physical sciences and engineering.
- SGI's **OpenInventor**, built atop OpenGL, SGI's 2D and 3D graphics library. This standard has a lot of history and support; it is the front-runner.
- **TSIPP**, an extension of Tcl (the Tool Control Language, see **Release 1.0**, 2-94), developed by John Ellson.

The participants will vote soon; then some of them will develop browsers, compilers and authoring tools that conform to the newly anointed VRML 1.0 specification. Later, they expect to offer the specification to the re-

---

6 *Mailing lists are highly democratic discussion tools because they can include users with the lowest possible Internet connectivity: an e-mail gateway to the Internet. Any message sent to the list address is reflected to all participants, who can subscribe or unsubscribe as they see fit by sending messages to a separate administrative address.*

7 *The name VRML, coined by Dave Raggett of HP Labs in Bristol, UK, reflects its ties to HTML, the Web's HyperText Markup Language.*

cently formed World Wide Web Consortium, which includes MIT, CERN (where the Web was created) and the National Center for Supercomputing Applications (where Mosaic was developed). By that time, VRML will have been through considerable peer review, which should speed acceptance. Eventually the members of the VRML list may present it to the Internet Engineering Task Force, which sets "official" Internet standards.

The VRML 1.0 spec will not support real-time interaction between multiple avatars, but will focus on making possible a concatenated "space" designed and shared by many individuals. In principle, object attributes could change frequently (e.g., through database lookups at runtime), though it is not easy to have them change once they are loaded. The designers of VRML 1.0 are leaving room to include interactivity in 2.0.

Behlendorf has plenty of experience moderating lists: He started the SFRaves mailing list several years ago. He's an undergraduate student in Computer Science at UC Berkeley and Wired Magazine's Webmaster. At Wired, he's helping to create HotWired, the Magazine's online presence, which will likely use VRML when it's available.

Pesce and the Labyrinth Group

At the Web conference, Pesce, president of the Labyrinth Group, described his proof-of-concept scene-description language, also called Labyrinth. The first two paragraphs of this issue are an embellished example of how Labyrinth could assemble and render scenes. Designers use it to define objects, scenes and camera positions, as well as other elements such as texture maps, lights and sounds. Each element might be stored on a different server; they are assembled dynamically when the scene is invoked.

Pesce displayed Labyrinth integrated with the World Wide Web at SIGGRAPH this month, in a special section called SIGKIDS. In collaboration with the US Holocaust Memorial Museum and Husky Labs (which manages the Museum's Web site), he created a virtual exhibit called Daniel's Story. The virtual exhibit recreates some of the rooms in the Holocaust Memorial. "Visitors" can move around in the rooms, walk around the objects in them (beds, desks, chairs, rugs, paintings) and read Daniel's diary pages, which describe the setting and the events.[8]

When visitors select a diary page, for example, they travel across the Web to typical Web pages. Conversely, when someone first invokes Daniel's Story from a Mosaic Web browser, MIME[9] automatically launches Labyrinth, which assembles the files, parses their contents and presents them. That makes it easy to move from Webspace to cyberspace and back.

---

8 *To view the starting page, point your graphical Web browser (e.g., Mosaic, MacWeb) to <http://www.butterfly.net/lab/daniel/sign.html>.*

9 *For more on MIME, the Multipurpose Internet Messaging Extensions, see Release 1.0, 2-94. MIME intercepts and dispatches requests for viewers, interpreters or for further processing that may accompany e-mail messages or Web network requests.*

Origins, destinations and the Cyberspace Protocol

Pesce created Labyrinth more to get people charged up about the possibilities of VR on the Web than to become the universal standard for VR on the Internet. He expects one of the submissions listed above to become the core of VRML 1.0, most likely OpenInventor, and will support whatever wins with public-domain browsers by the end of 1994. Now he funds operations from short-term projects. Interest in distributed virtual environments is heating up; Pesce is in discussion with companies about developing information services, training environments, commercial spaces and games.

Pesce has spent 15 years in networking and data communications, first as an engineer at Shiva, where he developed dial-up networking systems for Macs and PCs. William Gibson's novel **Neuromancer** triggered Pesce's longstanding interest in virtual environments and galvanized him into action. In 1991, Pesce started a company called Ono-Sendai to develop the software and hardware necessary for widely distributed virtual worlds (yes, Ono-Sendai is named after the vendor that makes the "deck" with which the protagonist in **Neuromancer** "jacks in" to cyberspace). He left Ono-Sendai in 1993 to start the Labyrinth Group.

Pesce believes that as more stuff is put into cyberspace, people will need some way to map and agree on a unified space -- a universal coordinate system that spans the entire Internet and provides conventions of scale and place. Gopher-style hierarchies and Web-style link traversal aren't helpful enough: Users get lost quickly and don't have enough meaningful context or framework.

He believes the underlying representation should be Cartesian, with some added features. Pesce has created a method for addressing and navigating virtual spaces called the Cyberspace Protocol, which he wants to offer to the Web and broader Internet communities. People can have their own mappings, says Pesce, but somewhere there should be a single 3D space with agreed-upon coordinates.

COMING SOON

- *The future of HTML and SGML.*
- *What's a zine?*
- *Smart communications software.*
- *Advertising online.*
- *Software for education.*
- *Intelligent drawing and authoring tools.*
- *And much more... (If you know of any good examples of the categories listed above, please let us know.)*

RESOURCES & PHONE NUMBERS

Ann Hardy, Dean Tribble, **Agorics**, (415) 941-8224; fax, (415) 941-8225; <annh@netcom.com>, <tribble@netcom.com>
Christopher Alexander, **Center for Environmental Structure**, (510) 841-6166; fax, (510) 841-8668
Carol Peters, Tom Melcher, **daVinci Time & Space**, (415) 525-2880; fax, (415) 525-2882; <melcher@dvts.com>
Chip Morningstar, **Electric Communities**, (415) 856-1130; fax, (415) 941-8225; <chip@netcom.com>
Mark Pesce, **Labyrinth Group**, (415) 621-1981; fax, (415) 474-1976; <mpesce@netcom.com>
Brian Behlendorf, **Wired Magazine**, (415) 904-0660; fax, (415) 904-0669; <brian@wired.com>
John Seely Brown, **Xerox PARC**, (415) 812-4341; fax, (415) 812-4037; <jsb.parc@xerox.com>
Pavel Curtis, **Xerox PARC**, (415) 812-4455; fax, (415) 812-4471; <pavel@parc.xerox.com>

For further reading:

Cyberspace: First Steps, edited by Michael Benedikt. MIT Press, 1991.
Christopher Alexander, **The Timeless Way of Building**, **A Pattern Language** and **The Oregon Experiment**. New York: Oxford University Press, 1979, 1977 and 1975.
Jean Lave and Etienne Wenger, **Situated Learning: Legitimate Peripheral Participation**. New York: Cambridge University Press, 1991.

Release 1.0 is published monthly, except for this combined July/August issue, by EDventure Holdings, 104 Fifth Ave., New York, NY 10011-6901; (212) 924-8800; fax, (212) 924-0240. It covers PCs, software, computer-telephone integration, groupware, text management, connectivity, messaging, wireless communications, artificial intelligence, intellectual property law and other unpredictable topics. A companion publication, Rel-EAST, is an information bulletin on emerging technology markets in Central Europe and the former Soviet units. Editor: Esther Dyson <esther@edventure.com>; publisher: Daphne Kis <daphne@edventure.com>; managing editor: Jerry Michalski <jerry@edventure.com>; circulation & fulfillment manager: Robyn Sturm <robyn@edventure.com>; executive assistant: Christina Koukkos; editorial & marketing communications consultant: William M. Kutik <kutik@edventure.com>.
Copyright 1994, EDventure Holdings Inc. All rights reserved. No material in this publication may be reproduced without written permission; however, we gladly arrange for reprints or bulk purchases. *Subscriptions cost $495 per year, $575 overseas.*

RELEASE 1.0 CALENDAR

July 31-Aug 4	**Artificial Intelligence conference** - Seattle. Sponsor: AAAI. Call Annette Eldredge, (415) 328-3123; fax, (415) 321-4457.
Aug 2-5	**MacWorld Expo** - Boston. Sponsor: MacWorld Magazine. Call David Eliot, (800) 945-3318; fax, (617) 361-9074.
August 4-6	**First annual conference on the pattern languages of programs** - Monticello, IL. Sponsor: University of Illinois, Department of Computer Science. Call Ralph Johnson, (217) 244-0093; fax, (217) 333-3501.
August 8-12	**GroupWare & Workflow '94** - San Jose. Sponsored by The Conference Group. Call David Coleman, (415) 282-9151; fax, (415) 550-8556.
August 10-11	**Transaction Switching and Transport Services Industry Forum** - Chicago. Sponsor: TSTS Industry Board. Call Cynthia Lucenius, (201) 740-6468; fax, (201) 740-6887.
August 14-17	**Software Marketing Perspectives '94** - San Francisco. Sponsor: Software Marketing Journal. Call Stuart Rauch, (415) 989-8765; fax, (415) 295-7766.
August 17-21	**@Online Networking Exposition and BBS Convention (ONE BBScon)** - Atlanta. Sponsor: ONE, Inc. Call Peg Coniglio, (303) 693-5253; fax, (303) 693-5518.
August 23-24	**Cyberspace and the American Dream** - Atlanta. Sponsor: Progress and Freedom Foundation. With Esther Dyson, George Gilder, Alvin Toffler, others. Call Bethany Noble (202) 484-2312; fax, (202) 484-9326; email pff@aol.com.
August 23-26	**Virtual reality software & technology** - Singapore. Sponsor: ACM and the Institute of Systems Science, Singapore. Call Gurminder Singh, (65) 772-3651; fax, (65) 774-4998.
Aug 29-Sept 1	**Client/Server East Conference & Exposition** - Washington, D.C. Sponsors: CMP Publications, Software Productivity Group, Sybase. Call Jonathan Varman, (516) 733-6827; fax, (516) 733-6753.
September 7-8	**Breaking the Barriers to the National Information Infrastructure** - Washington, D.C. Sponsor: Council on Competitiveness and the Information Infrastructure Task Force (IITF). Call Gurminder Singh, (65) 772-3651; fax, (65) 774-4998.
September 8-10	***ETRE** - Barcelona. Euro-Glitterati. Sponsored by Dasar. Call Alex Vieux, (415) 321-5544; fax, (415) 321-5597.
September 25-28	****East-West High-Tech Forum** - Bratislava, Slovakia. The fifth annual sponsored by EDventure Holdings. "Competition or Cooperation: Finding the right mix in a maturing market." Call Daphne Kis, (212) 924-8800; fax, (212) 924-0240; MCI 511-3763.

** Events Esther plans to attend.*
@ Events Jerry plans to attend.

Lack of a symbol is no indication of lack of merit.
Please let us know about other events we should include. -- Christina Koukkos

RELEASE 1.0

SUBSCRIPTION FORM

Please enter my subscription to **Release 1.0** at the rate of $495 per year in the U.S. and Canada. Overseas subscriptions are $575, airmail postage included. Payment must be enclosed. Multiple-copy rates available on request. Satisfaction guaranteed or your money back.

Name _____

Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Telephone _____

❑ Check enclosed.

❑ Charge my

❑ American Express ❑ MasterCard ❑ Visa

Card Number _____ Expiration Date _____

Name on Card _____ Signature _____

8-94

❑ Please send me information on your multiple-copy rate.

Please fill in the information above and send to:

EDVENTURE HOLDINGS INC.
104 FIFTH AVENUE, 20TH FLOOR
NEW YORK, NY 10011

If you have any questions, please call us at 1 (212) 924-8800;
Fax 1 (212) 924-0240; e-mail MCI:511-3705;
Internet 0005113705@mcimail.com.

Daphne Kis
Publisher